

PATENT

10830.0100.NPUS00

APPLICATION FOR UNITED STATES LETTERS PATENT

for

**MULTI-THREADED WRITE INTERFACE AND
METHODS FOR INCREASING THE SINGLE FILE READ AND WRITE
THROUGHPUT OF A FILE SERVER**

By

Sachin Mullick

Jiannan Zheng

Xiaoye Jiang

Sorin Faibish

Peter Bixby

Express Mail Mailing Label No. EV 318618485 US

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to file servers and data processing networks. The present invention more particularly relates to a file server permitting concurrent writes from multiple clients to the same file. The invention specifically relates to increasing the single file write throughput of such a file server.

2. Description of the Related Art

Network data storage is most economically provided by an array of low-cost disk drives integrated with a large semiconductor cache memory. A number of data mover computers are used to interface the cached disk array to the network. The data mover computers perform file locking management and mapping of the network files to logical block addresses of storage in the cached disk array, and move data between network clients and the storage in the cached disk array.

Data consistency problems may arise if multiple clients or processes have concurrent access to read-write files. Typically write synchronization and file locking have been used to ensure data consistency. For example, the data write path for a file has been serialized by holding an exclusive lock on the file for the entire duration of creating a list of data buffers to be written to disk, allocating the actual on-disk storage, and writing to storage synchronously. Unfortunately, these methods involve considerable access delays due to contention for locks not only on the files but also on the file directories and a log used when committing data to storage. In order to reduce these delays, a file server may permit asynchronous writes in accordance with version 3 of the Network File System (NFS) protocol. See, for example, Vahalia et al. U.S. Patent

1 5,893,140 issued April 6, 1999, entitled "File Server Having a File System Cache and
2 Protocol for Truly Safe Asynchronous Writes," incorporated herein by reference. More
3 recently, byte range locking to a file has been proposed in version 4 of the NFS protocol.
4 (See NFS Version 3 Protocol Specification, RFC 1813, Sun Microsystems, Inc., June
5 1995, incorporated herein by reference, and NFS Version 4 Protocol Specification, RFC
6 3530, Sun Microsystems, Inc., April 2003, incorporated herein by reference.)

7 Asynchronous writes and range locking alone will not eliminate access delays due
8 to contention during allocation and commitment of file metadata. A Unix-based file in
9 particular contains considerable metadata in the inode for the file and in indirect blocks of
10 the file. The inode, for example, contains the date of creation, date of access, file name,
11 and location of the data blocks used by the file in bitmap format. The NFS protocol
12 specifies how this metadata must be managed. In order to comply with the NFS protocol,
13 each time a write operation occurs, access to the file is not allowed until the metadata is
14 updated on disk, both for read and write operations. In a network environment, multiple
15 clients may issue simultaneous writes to the same large file such as a database, resulting
16 in considerable access delay during allocation and commitment of file metadata.

17

18 **SUMMARY OF THE INVENTION**

19 In accordance with one aspect of the present invention, there is provided a method
20 of operating a network file server for providing clients with concurrent write access to a
21 file. The method includes the network file server responding to a concurrent write
22 request from a client by obtaining a lock for the file, and then preallocating a metadata
23 block for the file, and then releasing the lock for the file, and then asynchronously writing

1 to the file, and then obtaining the lock for the file, and then committing the metadata
2 block to the file, and then releasing the lock for the file.

3 In accordance with another aspect, the invention provides a method of operating a
4 network file server for providing clients with concurrent write access to a file. The
5 method includes the network file server responding to a concurrent write request from a
6 client by preallocating a block for the file, and then asynchronously writing to the file,
7 and then committing the block to the file. The asynchronous writing to the file includes a
8 partial write to a new block that has been copied at least in part from an original block of
9 the file. The method further includes checking a partial block conflict queue for a
10 conflict with a concurrent write to the new block, and upon finding an indication of a
11 conflict with a concurrent write to the new block, waiting until resolution of the conflict
12 with the concurrent write to the new block, and then performing the partial write to the
13 new block.

14 In accordance with another aspect, the invention provides a method of operating a
15 network file server for providing clients with concurrent write access to a file. The
16 method includes the network file server responding to a concurrent write request from a
17 client by preallocating a metadata block for the file, and then asynchronously writing to
18 the file, and then committing the metadata block to the file. The method further includes
19 gathering together preallocated metadata blocks for a plurality of client write requests to
20 the file, and committing together the preallocated metadata blocks for the plurality of
21 client write requests to the file by obtaining a lock for the file, committing the gathered
22 preallocated metadata blocks for the plurality of client write requests to the file, and then
23 releasing the lock for the file.

1 In accordance with yet another aspect, the invention provides a method of
2 operating a network file server for providing clients with concurrent write access to a file.
3 The method includes the network file server responding to a concurrent write request
4 from a client by executing a write thread. Execution of the write thread includes
5 obtaining an allocation mutex for the file, and then preallocating new metadata blocks
6 that need to be allocated for writing to the file, and then releasing the allocation mutex for
7 the file, and then issuing asynchronous write requests for writing to the file, waiting for
8 callbacks indicating completion of the asynchronous write requests,
9 obtaining the allocation mutex for the file, and then committing the preallocated metadata
10 blocks, and then releasing the allocation mutex for the file.

11 In accordance with another aspect, the invention provides a network file server.
12 The network file server includes storage for storing a file and at least one processor
13 coupled to the storage for providing clients with concurrent write access to the file. The
14 network file server is programmed for responding to a concurrent write request from a
15 client by obtaining a lock for the file, and then preallocating a metadata block for the file,
16 and then releasing the lock for the file, and then asynchronously writing to the file, and
17 then obtaining the lock for the file, and then committing the metadata block to the file,
18 and then releasing the lock for the file.

19 In accordance with another aspect, the invention provides a network file server.
20 The network file server includes storage for storing a file, and at least one processor
21 coupled to the storage for providing clients with concurrent write access to the file. The
22 network file server is programmed for responding to a concurrent write request from a
23 client by preallocating a block for the file, and then asynchronously writing to the file,

1 and then committing the block to the file. The network file server includes a partial block
2 conflict queue for indicating a concurrent write to a new block that is being copied at
3 least in part from an original block of the file. The network file server is programmed for
4 responding to a client request for a partial write to the new block by checking the partial
5 block conflict queue for a conflict, and upon finding an indication of a conflict, waiting
6 until resolution of the conflict with the concurrent write to the new block of the file, and
7 then performing the partial write to the new block of the file.

8 In accordance with another aspect, the invention provides a network file server.
9 The network file server includes storage for storing a file, and at least one processor
10 coupled to the storage for providing clients with concurrent write access to the file. The
11 network file server is programmed for responding to a concurrent write request from a
12 client by preallocating a metadata block for the file, and then asynchronously writing to
13 the file; and then committing the metadata block to the file. The network file server is
14 programmed for gathering together preallocated metadata blocks for a plurality of client
15 write requests to the file, and committing together the preallocated metadata blocks for
16 the plurality of client write requests to the file by obtaining a lock for the file, committing
17 the gathered preallocated metadata blocks for the plurality of client write requests to the
18 file, and then releasing the lock for the file.

19 In accordance with yet still another aspect, the invention provides a network file
20 server. The network file server includes storage for storing a file, and at least one
21 processor coupled to the storage for providing clients with concurrent write access to the
22 file. The network file server is programmed with a write thread for responding to a
23 concurrent write request from a client by obtaining an allocation mutex for the file, and

1 then preallocating new metadata blocks that need to be allocated for writing to the file,
2 and then releasing the allocation mutex for the file, and then issuing asynchronous write
3 requests for writing to the file, waiting for callbacks indicating completion of the
4 asynchronous write requests, and then obtaining the allocation mutex for the file, and
5 then committing the preallocated metadata blocks, and then releasing the allocation
6 mutex for the file.

7 In accordance with a final aspect, the invention provides a network file server.
8 The network file server includes storage for storing a file, and at least one processor
9 coupled to the storage for providing clients with concurrent write access to the file. The
10 network file server is programmed for responding to a concurrent write request from a
11 client by preallocating a block for writing to the file, asynchronously writing to the file,
12 and then committing the preallocated block. The network file server also includes an
13 uncached write interface, a file system cache, and a cached read-write interface. The
14 uncached write interface bypasses the file system cache for sector-aligned write
15 operations, and the network file server is programmed to invalidate cache blocks in the
16 file system cache including sectors being written to by the cached read-write interface.

17

18 **BRIEF DESCRIPTION OF THE DRAWINGS**

19 Other objects and advantages of the invention will become apparent upon reading
20 the following detailed description with reference to the accompanying drawings wherein:

21 FIG. 1 is a block diagram of a data processing system including multiple clients
22 and a network file server;

1 FIG. 2 is a block diagram showing further details of the network file server in the
2 data processing system of FIG. 1;

3 FIG. 3 is a block diagram of various read and write interfaces in a Unix-based file
4 system layer (UxFS) in the network file server of FIG. 2;

5 FIG. 4 shows various file system data structures associated with a file in the
6 network file server of FIG. 2;

7 FIGS. 5 and 6 comprise a flowchart of programming in the Common File System
8 (CFS) layer in the network file server for handling a write request from a client;

9 FIG. 7 is a timing diagram showing multiple read and write operations pipelined
10 into parallel streams in the Common File System (CFS) layer in the network file server
11 for handling concurrent write requests from a client;

12 FIG. 8 shows multiple processors for processing the pipelined read and write
13 operations in the network file server;

14 FIG. 9 is a flowchart of programming in the Common File System (CFS) layer in
15 the network file server for handling a read request from a client;

16 FIG. 10 is a flowchart of programming in the Common File System (CFS) layer
17 in the network file server for handling concurrent read and write requests from a client;

18 FIG. 11 is a flowchart of a write thread in the UxFS layer of the network file
19 server;

20 FIG. 12 is a more detailed flowchart of steps in the write thread for committing
21 preallocated metadata;

22 FIG. 13 is a block diagram of a partial block write during a copy-on-write
23 operation;

1 FIG. 14 is a block diagram of a read-write file as maintained by the UxFS layer;

2 FIG. 15 is a block diagram of the read-write file of FIG. 14 after creation of a

3 read-only snapshot copy of the read-write file;

4 FIG. 16 is a block diagram of the read-write file of FIG. 15 after a copy-on-write

5 operation upon a data block and two indirect blocks between the data block and the inode

6 of the read-write file;

7 FIG. 17 is a flowchart of steps in a write thread for performing the partial block

8 write operation of FIG. 13;

9 FIG. 18 shows a flowchart of steps in a write thread for allocating file blocks

10 when writing to a file having read-only snapshots;

11 FIG. 19 is a block diagram of a file version set including read-only and read-write

12 snapshot copies of a production file;

13 FIG. 20 is a flowchart of a procedure for creating a new production file;

14 FIG. 21 is a block diagram of a conventional inode of a file;

15 FIG. 22 is a block diagram of an inode in the file version set of FIG. 19;

16 FIG. 23 is a block diagram showing linkages between the inodes in the file

17 version set of FIG. 19;

18 FIG. 24 is a flowchart of a procedure for creating a read-only snapshot copy in the

19 file version set of FIG. 19;

20 FIG. 25 is a flowchart of a procedure for creating a read-write branch in the file

21 version set of FIG. 19;

22 FIG. 26 is a flowchart of a procedure for deleting a read-only version in the file

23 version set of FIG. 19;

1 FIGS. 27-28 comprise a flowchart of a procedure for reserving file blocks for
2 read-write files in order to ensure that the sharing of file blocks among the files in the
3 version set of FIG. 19 is not likely to result in a shortage of file blocks when writing to
4 the read-write files;

5 FIG. 29 is a state diagram for the process of restoring a production file with a
6 read-only version;

7 FIG. 30 is a flowchart of a procedure for preparing for the restoration of the
8 production file with a read-only version;

9 FIG. 31 is a flowchart of a procedure for aborting the restoration of the production
10 file with a read-only version;

11 FIG. 32 is a flowchart of a procedure for committing the restoration of the
12 production file with a read-only version;

13 FIG. 33 is a flowchart of a procedure for refreshing a read-only version; and

14 FIGS. 34 and 35 comprise a flowchart of a procedure for parsing a file name for a
15 file in the version set of FIG. 19.

16 While the invention is susceptible to various modifications and alternative forms,
17 specific embodiments thereof have been shown in the drawings and will be described in
18 detail. It should be understood, however, that it is not intended to limit the invention to
19 the particular forms shown, but on the contrary, the intention is to cover all
20 modifications, equivalents, and alternatives falling within the scope of the invention as
21 defined by the appended claims.

22

DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

I. The preferred network file server.

FIG. 1 shows an Internet Protocol (IP) network 20 including a network file server 21 and multiple clients 23, 24, 25. The network file server 21, for example, has multiple data mover computers 26, 27, 28 for moving data between the IP network 20 and a cached disk array 29. The network file server 21 also has a control station 30 connected via a dedicated dual-redundant data link 31 among the data movers for configuring the data movers and the cached disk array 29.

Further details regarding the network file server 21 are found in Vahalia et al., U.S. Patent 5,893,140, incorporated herein by reference, and Xu et al., U.S. Patent 6,324,581, issued Nov. 27, 2001, incorporated herein by reference. The network file server 21 is managed as a dedicated network appliance, integrated with popular network operating systems in a way, which, other than its superior performance, is transparent to the end user. The clustering of the data movers 26, 27, 28 as a front end to the cached disk array 29 provides parallelism and scalability. Each of the data movers 26, 27, 28 is a high-end commodity computer, providing the highest performance appropriate for a data mover at the lowest cost. The data mover computers 26, 27, 28 may communicate with the other network devices using standard file access protocols such as the Network File System (NFS) or the Common Internet File System (CIFS) protocols, but the data mover computers do not necessarily employ standard operating systems. For example, the network file server 21 is programmed with a Unix-based file system that has been adapted for rapid file access and streaming of data between the cached disk array 29 and the data network 20 by any one of the data mover computers 26, 27, 28.

1 FIG. 2 shows software modules in the data mover 26 introduced in FIG. 1. The
2 data mover 26 has a Network File System (NFS) module 41 for supporting
3 communication among the clients and data movers of FIG. 1 over the IP network 20
4 using the NFS file access protocol, and a Common Internet File System (CIFS) module
5 42 for supporting communication over the IP network using the CIFS file access
6 protocol. The NFS module 41 and the CIFS module 42 are layered over a Common File
7 System (CFS) module 43, and the CFS module is layered over a Universal File System
8 (UxFS) module 44. The UxFS module supports a UNIX-based file system, and the CFS
9 module 43 provides higher-level functions common to NFS and CIFS.

10 The UxFS module accesses data organized into logical volumes defined by a
11 module 45. Each logical volume maps to contiguous logical storage addresses in the
12 cached disk array 29. The module 45 is layered over a SCSI driver 46 and a Fibre-
13 channel protocol (FCP) driver 47. The data mover 26 sends storage access requests
14 through a host bus adapter 48 using the SCSI protocol, the iSCSI protocol, or the Fibre-
15 Channel protocol, depending on the physical link between the data mover 26 and the
16 cached disk array 29.

17 A network interface card 49 in the data mover 26 receives IP data packets from
18 the IP network 20. A TCP/IP module 50 decodes data from the IP data packets for the
19 TCP connection and stores the data in message buffers 53. For example, the UxFS layer
20 44 writes data from the message buffers 53 to a file system 54 in the cached disk array
21 29. The UxFS layer 44 also reads data from the file system 54 or a file system cache 51
22 and copies the data into the message buffers 53 for transmission to the network clients 23,
23 24, 25.

1 To maintain the file system 54 in a consistent state during concurrent writes to a
2 file, the UxFS layer maintains file system data structures 52 in random access memory of
3 the data mover 26. To enable recovery of the file system 54 to a consistent state after a
4 system crash, the UxFS layer writes file metadata to a log 55 in the cached disk array
5 during the commit of certain write operations to the file system 54.

6

7 II. Concurrent read and write operations and the allocation of file system blocks.

8 FIG. 3 shows various read and write interfaces in the UxFS layer. These
9 interfaces include a cached read/write interface 61 for accessing the file system cache 51,
10 an uncached multi-threaded write interface 63, and an uncached read interface 64.

11 The cached read/write interface 61 permits reads and writes to the file system
12 cache 51. If data to be accessed does not reside in the cache, it is staged from the file
13 system 54 to the file system cache 51. Data written to the file system cache 51 from the
14 cached read/write interface 61 is written down to the file system cache during a commit
15 operation. The file data is written down first, followed by writing of new file metadata
16 to the log 55 and then writing of the new metadata to the file system 54.

17 The uncached multi-threaded write interface 63 is used for sector-aligned writes
18 to the file system 54. Sectors of data (e.g., 512 byte blocks) are read from the message
19 buffers (53 in FIG. 2) and written directly to the cached disk array 29. For example, each
20 file block is sector aligned and is 8 K bytes in length. When a sector-aligned write
21 occurs, any cache blocks in the file system cache that include the sectors being written to
22 are invalidated. In effect, the uncached multi-threaded write interface 63 commits file
23 data when writing the file data to the file system 54 in storage. The uncached multi-

1 threaded write interface 63 allows multiple concurrent writes to the same file. If a sector-
2 aligned write changes metadata of a file such as file block allocations, then after the data
3 of the file has been written, the new metadata is written to the log 55, and then the new
4 metadata is written to the file system 54. The new metadata includes modifications to the
5 file's inode, any new or modified indirect blocks, and any modified quota reservation.

6 The uncached read interface 64 reads sectors of data directly from the file system
7 54 into the message buffers (53 in FIG. 2). For example, the read request must have a
8 sector aligned offset and specifies a sector count for the amount of data to be read. The
9 data can be read into multiple message buffers in one input/output operation so long as
10 the sectors to be read are in contiguous file system blocks.

11 Typically, the cached read/write interface 61 is used for reading data from read-
12 write files and from any read-only snapshot copies of the read-write files. The uncached
13 write interface 63 is used for sector-aligned writes to read-write files. If the writes are
14 not sector aligned, then the cached read-write interface 61 is used. The uncached read
15 interface 64 is used for sector-aligned reads when there is no advantage to retaining the
16 data in the file system cache 51; for example, when streaming data to a remote copy of a
17 file.

18 FIG. 4 shows various file system data structures 52 associated with a file. A
19 virtual inode (VNODE) 71 represents the file. The virtual inode 71 is linked to an
20 allocation mutex (mutually exclusive lock) 72, a partial block conflict queue 73, a partial
21 write wait queue 74, an input-output (I/O) list 75, a staging queue 76, and preallocation
22 block lists 77. When a file block is preallocated, it is reserved for use in the on-disk file
23 system 54. A preallocated file block can be linked into the in-memory file block

1 structure in the file system cache 51 as maintained by the UxFS layer 44, and later the
2 preallocated file block can become part of the on-disk file system 54 when the
3 preallocated file block is committed to storage. (An example of the file block structure is
4 shown in FIG. 14.) The write threads of the uncached multi-threaded write interface (63
5 in FIG. 3) use the allocation mutex 72 for serializing preallocation of file metadata blocks
6 and commitment of the preallocated metadata blocks. For a Unix-based file, the
7 preallocated metadata blocks include new indirect blocks, which are added to the file
8 when the file is extended. As described below with reference to FIGS. 15 to 16, one or
9 more new indirect blocks may also be added to a read-write file system when processing
10 a client request to write to a data block that is shared between the read-write file system
11 and a read-only snapshot copy of the read-write file system.

12 Preallocation of the file metadata blocks under control of the allocation mutex
13 prevents multiple writers from allocating the same metadata block. The actual data write
14 is done using asynchronous callbacks within the context of the thread, and does not hold
15 any locks. Since writing to the on-disk storage takes the majority of the time, the
16 preallocation method enhances concurrency, while maintaining data integrity.

17 The preallocation method allows concurrent writes to indirect blocks within the
18 same file. Multiple writers can write to the same indirect block tree concurrently without
19 improper replication of the indirect blocks. Two different indirect blocks will not be
20 allocated for replicating the same indirect block. The write threads use the partial block
21 conflict queue 73 and the partial write wait queue 74 to avoid conflict during partial
22 block write operations, as further described below with reference to FIG. 13.

1 The I/O list 75 maps the message buffers (53 in FIG. 2) to data blocks to be
2 written. The write threads use the I/O list 75 to implement byte range locking. The read
3 threads may also use the I/O for byte-range locking. The data blocks, for example, are
4 512 bytes in length providing sector-level granularity for the byte range locking.
5 Alternatively, the data block length is a multiple of the sector size.

6 In order to prevent the log (55 in FIG. 2) from becoming a bottleneck, the
7 preallocated metadata blocks for multiple write threads writing to the file at the same
8 time are committed together under the same logging lock. Committing more than one
9 allocation under one lock increases the throughput. For this purpose, a staging queue 76
10 is allocated and linked to the file virtual inode 71. Preallocation block lists 77 identify
11 the respective preallocated metadata blocks for the write threads writing to the file. The
12 staging queue 76 receives pointers to the preallocation block lists 77 of the write threads
13 waiting for the allocation mutex 72 of the file for commitment of their preallocated
14 metadata blocks. For example, the staging queue 76 is a conventional circular queue, or
15 the preallocation block lists 77 are linked together into a circular list to form the staging
16 queue. There can be multiple files, and each file can have a respective staging queue
17 waiting for commitment of the file's preallocation block lists. A wait list of staging
18 queues 78 identifies the staging queues waiting for service on a first-come, first-served
19 basis.

20 From a client's view, the write operation performed by a write thread in the
21 uncached write interface is a synchronous operation. The write thread does not return an
22 acknowledgement to the client until the write data has been written down to the file
23 system in storage, and the metadata allocation has been committed to storage.

1 FIGS. 5 and 6 show programming in the Common File System (CFS) layer in the
2 network file server for handling a write request from a client. In a first step 81, if the
3 uncached multi-threaded write interface (63 in FIG. 3) is not turned on for the file
4 system, then execution branches to step 82. For example, the uncached interface can be
5 turned on or off per file system as a mount-time option. In step 82, the CFS layer obtains
6 an exclusive lock upon the file, for example by acquiring the allocation mutex (72 in FIG.
7 4) for the file. Then in step 83, the CFS layer writes a specified number of bytes from the
8 source to the file, starting at a specified byte offset, using the cached read/write interface
9 (61 in FIG. 3). The source, for example, is one or more of the message buffers (53).
10 Then in step 84, the CFS layer releases the exclusive lock upon the file, and processing of
11 the write request is finished.

12 In step 81, if the uncached multi-threaded write interface is turned on for the file
13 system, then execution continues to step 85. In step 85, if the write data specified by the
14 write request is not sector aligned (or the data size is not in multiple sectors), then
15 execution branches to step 82. Otherwise, execution continues from step 85 to step 86.

16 In step 86, the CFS layer acquires a shared lock upon the file. The shared lock
17 prevents the CFS layer from obtaining an exclusive lock upon the file for a concurrent
18 write request (e.g., in step 82). However, as described below, the shared lock upon the
19 file does not prohibit write threads in the UxFS layer from acquiring the allocation mutex
20 (72 in FIG. 4) during the preallocation of metadata blocks or during the commitment of
21 the metadata blocks.

22 In step 87, the CFS layer checks the I/O list (75 in FIG. 4) for a conflict. If there
23 is a conflicting data block on the I/O list, then execution waits until the conflicting data

1 block is flushed out of the I/O list. For example, for serializing the writes with prior
2 reads and writes, write access to any blocks being accessed by prior in-progress reads or
3 writes is delayed until these blocks have been accessed by these prior in-progress reads or
4 writes. Moreover, in certain clustered systems in which direct data access to the file in
5 the data storage is shared with other servers or clients, execution may also wait in step 87
6 for range locks to be released by another server or client sharing direct access to the file.
7 After step 87, execution continues to step 88 in FIG. 6.

8 In step 88 of FIG. 6, the CFS layer writes the specified number of bytes from the
9 source to the file, starting at a specified sector offset, using the uncached multi-threaded
10 write interface (63 in FIG. 3). Then in step 89, the CFS layer invalidates any cached
11 entries for the file system blocks that have been written to in the file system cache (51 in
12 FIG. 3). The invalidation occurs after completion of any reads in progress to these file
13 system blocks. In step 90, the CFS layer releases the shared lock upon the file, and
14 processing of the write request is finished.

15 FIG. 7 shows that the parallel read and write architecture can be used to achieve
16 pipelining, since the data write stage does not involve any metadata interactions. The
17 read or write is divided into three steps, namely inode access for reads and writes and
18 preallocation for writes (S1), asynchronous read or write (S2), and inode access for reads
19 and writes and commit for writes (S3). The preallocation in stage S1 is achieved
20 synchronously, and the allocation mutex (72 in FIG. 4) prevents multiple preallocations
21 from occurring simultaneously for the same file.

22 Once the metadata preallocation stage is complete, the asynchronous write (S2) of
23 the data to disk can be handled independently of the metadata preallocation. The

1 asynchronous write (in stage S2) of the data to disk is the longest stage. With pipelining,
2 multiple asynchronous writes can be performed concurrently. This results in an increase
3 in the number of write operations that can be performed in a given time period.

4 The final commit of the allocations (in stage S3) is also achieved synchronously.
5 The allocation mutex (72 in FIG. 4) prevents preallocation for the same file from
6 occurring at the same time as a commit for the same file. However, multiple commits
7 (S3) for the same file may occur simultaneously by gathering the commit requests
8 together and committing them under the same allocation mutex.

9 As shown in FIG. 8, the read and write operations can be pipelined through
10 multiple processors. In this example, a multi-processor board 501 includes two Pentium
11 IV Xeon™ processor chips 502, 503. Each processor chip includes two logical central
12 processing units (CPU) 504, 505 and 506, 507 respectively. Each logical CPU consists
13 of a respective set of on-chip processor registers that share the functional units, input-
14 output ports and cache memory on the chip.

15 The processing of a multi-threaded application by the two logical processors on
16 the same physical processor is called “Hyper-Threading technology.” See, for example,
17 “Building Cutting-Edge Server Applications, Intel® Xeon™ Processor Family Features
18 the Intel NetBurst™ Microarchitecture with Hyper-Threading Technology,” Intel
19 Corporation, 2002, and Chapter 7, “Multiprocessor and Hyper-Threading Technology,”
20 in the Intel® Pentium™ 4 and Intel® Xeon™ Processor Optimization Reference Manual,
21 Order No. 248966-05, Intel Corporation, 2002.

22 For pipelined processing of the multiple concurrent read and write operations as
23 shown in FIG. 7, the metadata management for a file (stages S1 and S3) can be

1 performed by one logical processor, and the asynchronous reads and writes (stage S2) can
2 be performed by another logical processor. Each logical processor executes code threads
3 that are independent of the code threads executed by the other logical processors. For
4 example, as shown in FIG. 8, the first logical CPU 504 of the first processor chip 502
5 performs metadata management for files in a first file system (A:), and the first logical
6 CPU 506 of the second processor chip 503 performs metadata management for files in a
7 second file system (B:). The second logical CPU 504, 507 in each processor chip 502,
8 503 performs asynchronous write operations. For example, the second logical CPU 505
9 of the first processor chip 502 performs asynchronous read and write operations for the
10 first file system (A:), and if the second logical CPU 505 has free processing time, then the
11 second logical CPU 505 of the first processor chip 502 performs asynchronous read and
12 write operations for the second file system (B:). In a similar fashion, the second logical
13 CPU 507 of the second processor chip 502 performs asynchronous read and write
14 operations for the second file system (B:), and if the second logical CPU 507 has free
15 processing time, then the second logical CPU 505 of the first processor chip 502
16 performs asynchronous read and write operations for the first file system (A:).

17 In general, data read requests can be pipelined along with the write requests, and
18 separate processing units can service data read and write requests generated by a primary
19 processor that handles metadata management for each file. The primary processor can
20 hand over an input/output list to a separate secondary processing unit that will then go
21 through the input/output list to perform the transfer of data between the message buffers
22 and cache or disk. For a write operation, the separate secondary processing unit will take
23 the data from the network packets, write it to specified disk locations as requested by the

1 primary, and complete the data write to the disk from the network packets. The pipeline
2 architecture allows a primary processor to do the next write metadata preallocation while
3 other secondary processors are still writing data to disk.

4 When a write I/O request arrives at a primary processor or thread, the request is
5 analyzed and if there are any associated metadata operations (stage S1 in FIG. 7), and the
6 associated metadata operations are executed by the primary processor while the block
7 write I/O is pipelined to another separate secondary processing unit. The secondary
8 processing unit will pipeline multiple block write I/Os (stage S2 in FIG. 7), and will
9 commit the write data to the disk independently of the metadata operation. At the end of
10 the write data commit process, the metadata is committed (stage S3 in FIG. 7) to disk as
11 well. The primary processor is freed to perform additional metadata management
12 operations while the secondary processing unit writes the I/O data to the disk.

13 There could be a pool of secondary logical processing units that execute the write
14 tasks and they can be allocated for additional processing tasks by the primary processor.
15 Some tasks are executed only by the primary processor. For example, a primary
16 processor is allocated to a file system when the data mover is rebooted. The processing
17 of the pipeline is based on the fact that the writes are uncached, and once an
18 asynchronous write is issued to a secondary processor, there is no contingency or locking
19 to the file. If there are any contingencies, then they are solved by the primary processor
20 before the write is issued to the secondary processor.

21 FIG. 9 shows programming in the CFS layer in the network file server for
22 handling a read request from a client concurrent with handling a write request to the same

1 file. For clarity, FIG. 9 omits certain steps for handling read-write interactions. These
2 steps are show in FIG. 10 and further described below.

3 In a first step 91 of FIG. 9, the CFS layer obtains a shared lock upon the file for
4 the read request. In addition, for serializing the read with prior writes, the I/O list for the
5 file (75 in FIG. 4) can be accessed, and read access (in steps 93 or 97) to any blocks
6 being written to by prior in-progress writes is delayed until these blocks have been
7 written to by these prior in-progress writes. In step 92, execution continues to step 93 if
8 the data requested by the read request is found in the file system cache (51 in FIG. 2). In
9 step 93, the data requested by the read request is read from the cache. In step 94, the data
10 read from the cache is added to source message buffers (53 in FIG. 2). In step 95, the
11 shared lock upon the file is released for the write request, and the handling of the read
12 request by the CFS layer is done.

13 In step 92, in the requested data is not found in the cache, then execution branches
14 to step 96. In step 96, the last committed version of the inode for the file is accessed to
15 perform a search for the data block on disk containing the requested data. In step 97, the
16 requested data is read from the data block on disk. In step 98, the data read from the data
17 block on disk is added to the cache. After step 98, execution continues to step 94 to add
18 the data to the source message buffers.

19 FIG. 10 shows the behavior of the server when there are read-write interactions
20 during concurrent access of multiple I/O threads to a single file. In particular, the steps in
21 FIG. 10 occur when a read I/O request accesses blocks to which there is a concurrent
22 ongoing write. Each read and write must obtain a shared lock upon the file. A read

1 acquires the shared lock upon the file in step 91, and a write acquires the shared lock
2 upon the file in step 86.

3 The file system cache (51 in FIG. 2) maintains an index or block map including,
4 for each file system block, an I/O in progress (IOP) flag indicating whether or not a read
5 to the file system block on disk is in progress. During a cached read, if the block map
6 indicates that the block of data to be read is not found in the file system, then execution
7 branches from step 92 to step 510. In step 510, if the IOP flag is set, then execution
8 continues to step 511 to wait for the IOP flag to be cleared. Execution loops back from
9 step 511 to step 92.

10 If in step 510 the IOP flag is not set, then execution branches to step 512 to set the
11 IOP flag and to set a generation count for the block to a value of the present read of the
12 block from disk, and to start the read of the disk. The read of the disk is performed in
13 step 96 to get the committed mapping from the inode, and in step 97 to read data from the
14 block on disk. Once the data from the disk is obtained, the IOP flag and the generation
15 count are checked in step 513. If the IOP flag is set and the generation count is the same
16 as it was in step 512 for the read operation, then in step 98 the data is added to the read
17 cache. After step 98, execution continues to step 94.

18 It is possible that in step 513, the IOP flag will be cleared, or the generation count
19 may be different. The IOP flag can be cleared by a concurrent write operation. For
20 example, after a shared lock upon the file is obtained in step 86 for a concurrent write to
21 the file, data is written in step 515 from message buffers to disk. After step 515, in step
22 516, any cache data for the data block is invalidated, and any IOP flag for the data block
23 is cleared. After step 516, execution continues to step 95 to release the shared lock upon

1 the file. It is possible for a concurrent read to this file block to begin just after the IOP
2 flag for the block is cleared in step 516 but before a previous read has reached step 513.
3 In this case, the previous read will find that the IOP flag is set in step 513 but the
4 generation count will have changed, so that step 98 of adding the data to the cache will be
5 skipped. Step 98 is skipped under these circumstances because the data is current for this
6 previous read operation but stale for subsequent read operations.

7 In short, a write request is serviced by finding partial blocks and creating a partial
8 block list, preallocating metadata blocks for the range of block numbers in the inode that
9 is being written, issuing asynchronous write requests, waiting for completion of the
10 asynchronous write requests, getting a block commit lock, committing the preallocated
11 metadata blocks for the range written to in the inode, releasing the block commit lock,
12 starting asynchronous writes for conflict I/Os, finding the range of blocks in the file
13 system cache to be invalidated, invalidating the file system cache blocks for the block
14 range being committed, if there are active readers, marking the cache range as stale data
15 (for example, via the IOP flag and generation count mechanism shown in FIG. 10).

16 In short, a read request serviced by finding the range of file blocks to be read, and
17 if the blocks are in cache, then reading the data from the cache, and if not, then getting
18 the block commit lock, getting the committed mapping from the inode for the range of
19 blocks to be read, releasing the block commit lock (i.e., the allocation mutex), reading
20 data from disk to the buffer cache and source, and if there are stale blocks in the block
21 range (because a write to the blocks occurred during the read), then invalidating the stale
22 blocks.

1 During a read, IOP flags and generation counts can be used to identify stale
2 blocks. For example, when looking up to see whether data to be read is in cache, missing
3 blocks are marked as IOP (IO in Progress) and the generation count is set to a value
4 associated with this read, and then a read will be started. After completing any reads
5 necessary the blocks that were previously marked as IOP are cleared in one of the
6 following ways: (1) if the slot is cleared, then it's been purged and the just completed
7 read should not be entered and cached; (2) if it's marked as IOP then the generation count
8 is checked: if the generation count is the same as set for this read then data for this read is
9 cached in the slot; otherwise, the data for this read is not cached in this slot but otherwise
10 it can be used to satisfy the read request. During the read process, any concurrent writes
11 are simply allowed to proceed. At the end of the write, the entire range of blocks
12 written are invalidated in the cache. If a cache slot in the range is empty, then it is
13 ignored; otherwise, if the slot had a hint then the slot is cleared, and if the slot was IOP
14 then the IOP flag is cleared and any waiting reads (in step 511 of FIG. 10) are awoken
15 and allowed to proceed.

16 Servicing of the concurrent read and write requests as described above prevents
17 writes to a file from being blocked. There is, however, still an issue of concurrent reads
18 and writes to the same blocks in the same file. For some applications, it is desirable to
19 serialize these reads and writes in so that the data returned by a read operation will
20 indicate that the writes are atomic operations. For example, if each of two write
21 operations write to the same two blocks, a read should not return a first block from the
22 first write operation and a second block from the second write operation. This problem
23 can be solved by looking for conflicting blocks for prior in-progress reads and writes

1 before issuing an asynchronous write operation and by looking for conflicting blocks for
2 prior in-progress writes before issuing an asynchronous read operation, and if a conflict is
3 found, waiting for these prior in-progress conflicting operations to complete before the
4 asynchronous write operation or read operation is issued. This can be done by inspection
5 of the block ranges for prior in-progress writes in the I/O list 75 in FIG. 4. For reads, this
6 would be done in step 91 of FIG. 9. In addition, a read could immediately access non-
7 conflicting blocks in the cache, without waiting for the prior-in-progress writes to
8 complete.

9 Instead of using the I/O list to serialize reads and writes to the same file blocks, a
10 “Write In Progress” (WIP) flag could be added to the file system cache block map. In
11 effect, the WIP flag would be a write lock at the file block level of granularity. Before
12 issuing an asynchronous write operation, during the preallocation stage (S1), the primary
13 processor would set the WIP flags for the file system blocks being written to, unless a
14 WIP flag would already be set, in which case, the write operation would need to wait for
15 completion of the prior conflicting write. The WIP flags would be reset in the
16 asynchronous write stage (S2) after writing to each block. Subsequent writes that
17 encountered a set WIP flag within it’s own block range would be required to wait before
18 writing to each block. Likewise attempts to read that encounter a set WIP flag would
19 need to wait until the WIP flag is reset by completion of the conflicting write. If a read
20 operation is accessing blocks being written to by prior in-progress writes, then the read
21 operation should not access these blocks until after they have been written to by the prior
22 in-progress writes. For example, in step 91 of FIG. 9, the I/O list (75 in FIG. 4) or the
23 WIP flags can be accessed to determine the conflicting blocks, before attempting to

1 access these blocks in cache. However, the cache can be accessed immediately for
2 blocks that are not being written to by prior, in-progress writes.

3 FIG. 11 shows a flowchart of a write thread in the UxFS layer (44 in FIG. 2). In a
4 first step 101, the write thread gets the allocation mutex (72 in FIG. 4) for the file. Then
5 in step 102, the write thread preallocates metadata blocks for the block range being
6 written to the file. In step 103, the write thread releases the allocation mutex for the file.

7 In step 104, the write thread issues asynchronous write requests for writing to
8 blocks of the file. For example, a list of callbacks is created. There is one callback for
9 each asynchronous write request consisting of up to 64 K bytes of data from one or more
10 contiguous file system blocks. An I/O list is created for each callback. The
11 asynchronous write requests are issued asynchronously, so multiple asynchronous writes
12 may be in progress concurrently. In step 105, the write thread waits for the asynchronous
13 write requests to complete.

14 In step 106, the write thread gets the allocation mutex for the file. In step 107, the
15 write thread commits the preallocated metadata blocks to the file system in storage. The
16 new metadata for the file including the preallocated metadata blocks is committed by
17 being written to the log (55 in FIG. 3). File system metadata such as the file modification
18 time, however, is not committed in step 107 and is not logged. Instead, file system
19 metadata such as the file modification time is updated at a file system sync time during
20 the flushing of file system inodes. Finally, in step 108, the write thread releases the
21 allocation mutex for the file. This method of preallocating and committing metadata
22 blocks does not need any locking or metadata transactions for re-writing to allocated
23 blocks.

1 FIG. 12 is a more detailed flowchart of steps in the write thread for committing
2 the preallocated metadata. In a first step 111, if there is not a previous commit in
3 progress, then execution continues to step 112. In step 112, the thread gets the allocation
4 mutex for the file. Then in step 113, the thread writes new metadata (identified by the
5 thread's preallocation list) to the log in storage. In step 114, the thread writes the new
6 metadata (identified by the thread's preallocation list) to the file system in storage. In
7 step 115, the thread releases the allocation mutex for the file. Finally, in step 116, the
8 thread returns an acknowledgement of the write operation.

9 In step 111, if there was a previous commit in progress, then the thread inserts a
10 pointer to the threads' preallocation list onto the tail of the staging queue for the file. If
11 the staging queue was empty, then the staging queue is put on the wait list of staging
12 queues (78 in FIG. 4). The thread is suspended, waiting for a callback from servicing of
13 the staging queue. In step 118, the metadata identified by the thread's preallocation list is
14 committed when the staging queue is serviced. The staging queue is serviced by
15 obtaining the allocation mutex for the file, writing the new metadata for all of the
16 preallocation lists on the staging queue to the log in storage, then writing this new
17 metadata to the file system in storage, and then releasing the allocation mutex for the file.
18 Once servicing of the staging queue has committed the new metadata for the thread's
19 preallocation list, execution of the thread is resumed in step 116 to return an
20 acknowledgement of the write operation. After step 116, the thread is finished with the
21 write operation.

22 FIG. 13 is a block diagram of a partial block write during a copy-on-write
23 operation. Such an operation involves copying a portion of the data from an original file

1 system block 121 to a newly allocated file system block 123, and writing a new partial
2 block of data 122 to the newly allocated file system block. The portion of the data from
3 the original file system block becomes merged with the new partial block of data 122. If
4 the new partial block of data is sector aligned, then the partial block write can be
5 performed by the uncached multi-threaded write interface (63 in FIG. 3). Otherwise, if
6 the new partial block of data were not sector aligned, then the partial block write would
7 be performed by the cached read/write interface (61 in FIG. 3).

8 The copy-on-write operation may frequently occur in a file system including one
9 or more read-only file snapshot copies of a read-write file. Such a file system is
10 described in Chutani, Sailesh, et al., "The Episode File System," Carnegie Mellon
11 University IT Center, Pittsburgh, PA, June 1991, incorporated herein by reference. Each
12 read-only snapshot copy is the state of the read-write file at a respective point in time.
13 Read-only snapshot copies can be used for on-line data backup and data mining tasks.

14 In a copy-on-write file versioning method, the read-only snapshot copy initially
15 includes only a copy of the inode of the original file. Therefore the read-only snapshot
16 copy initially shares all of the data blocks as well as any indirect blocks of the original
17 file. When the original file is modified, new blocks are allocated and linked to the
18 original file inode to save the new data, and the original data blocks are retained and
19 linked to the inode of the read-only snapshot copy. The result is that disk space is saved
20 by only saving the difference between two consecutive snapshot copies. This process is
21 shown in FIGS. 13, 14, and 15.

22 FIG. 14 shows a read-write file as maintained by the UxFS layer. The file has a
23 hierarchical organization, depicted as an inverted tree. The file includes a read-write

1 inode 131, a data block 132 and an indirect block 133 linked to the read-write inode, a
2 data block 134 and an indirect block 135 linked to the indirect block 133, and data blocks
3 136 and 137 linked to the indirect block 135.

4 When a read-only snapshot copy of a read-write file is created, a new inode for
5 the read-only snapshot copy is allocated. The read-write file inode and file handle remain
6 the same. After allocation of the new inode, the read-write file is locked and the new
7 inode is populated from the contents of the read-write file inode. Then the read-write file
8 inode itself is modified, the transaction is committed, and the lock on the read-write file
9 is released.

10 The allocation of blocks during the copy-on-write to the read-write file raises the
11 possibility of the supply of free storage being used up after writing to a small fraction of
12 the blocks of the read-write file. To eliminate this possibility, the read-write file can be
13 provided with a "persistent reservation" mechanism so that the creation of a read-only
14 snapshot copy will fail unless there can be reserved a number of free storage blocks equal
15 to the number of blocks that become shared between the read-only snapshot copy and the
16 read-write file. The number of reserved blocks can be maintained as an attribute of the
17 file. The number of reserved blocks for a read-only file can be incremented as blocks
18 become shared with a read-only snapshot copy, and decremented as blocks are allocated
19 during the writes to the read-write file.

20 FIG. 15 shows the read-write file of FIG. 14 after creation of a read-only snapshot
21 copy of the read-write file. The read-only inode 138 is a copy of the read-write inode
22 131. The read-write inode 131 has been modified to indicate that the data block 132 and
23 the indirect block 133 are shared with a read-only snapshot copy. For example, in the

1 read-write inode 131, the most significant bit in each of the pointers to data block 132
2 and the indirect block 133 have been set to indicate that the pointers point to blocks that
3 are shared with the read-write file. (The links represented by such pointers to shared
4 blocks are indicated by dotted lines in FIGS. 15 and 16.) Also, by inheritance, any and
5 all of the descendants of a shared block are also shared blocks. Routines in the UxFS
6 layer that use the pointers to locate the pointed-to file system blocks simply mask out the
7 most significant to determine the block addresses.

8 In general, for the case in which there are multiple versions of a file sharing file
9 blocks, when a file block is shared, it is desirable to designate the oldest snapshot copy
10 sharing the block to be the owner of the block, and any other files to be non-owners of
11 the block. A pointer in a non-shared block pointing to a shared block will have its most
12 significant bit set if the block is not owned by the owner of the non-shared block, and will
13 have its most significant bit clear if the block is owned by the owner of the non-shared
14 block.

15 When writing to a specified sector of a file, a search of the file block hierarchy is
16 done starting with the read-write inode, in order to find the file block containing the
17 specified sector. Upon finding a pointer indicating that the pointed-to block is shared, the
18 pointed-to block and its descendants are noted as “copy on write” blocks. If the specified
19 sector is found in a “copy on write” block, then a new file block is allocated.

20 In practice, multiple write threads are executed concurrently, so that more than
21 one concurrent write thread could determine a need to preallocate the same new file
22 block. The allocation mutex is used to serialize the allocation process so more than one
23 preallocation of a new file block does not occur. For example, once the write thread has

1 obtained the allocation mutex, the write thread then determines whether a new block is
2 needed, and if so, then the write thread preallocates the new block. The write thread may
3 obtain the allocation mutex, allocate multiple new blocks in this fashion, and then release
4 the allocation mutex. For example, to write to a data block of a file, when the write
5 thread finds a shared block on the path in the file hierarchy down to the data block of the
6 file, the write thread obtains the allocation mutex, and then allocates all the shared blocks
7 that it then finds down the path in the file hierarchy down to and including the data block,
8 and then release the allocation mutex.

9 Once a new file block has been allocated, a partial block write to the new file
10 block is performed, unless the write operation writes new data to the entire block. The
11 new file block is the same type (direct or indirect) as the original “copy on write” file
12 block containing the specified sector. If the write operation writes new data to the entire
13 new file block, then no copy need be done and the new data is simply written into the
14 newly allocated block. (A partial write could be performed when the write operation
15 writes new data to the entire block, although this would not provide the best
16 performance.)

17 If the read-write inode or a block owned by the read-write file was a parent of the
18 original “copy on write” block, then the new file block becomes a child of the read-write
19 inode or the block owned by the read-write file. Otherwise, the new file block becomes
20 the child of a newly allocated indirect block. In particular, copies are made of all of the
21 “copy on write” indirect blocks that are descendants of the read-write inode and are also
22 predecessors of the original “copy on write” file block.

1 For example, assume that a write request specifies a sector found to be in the data
2 block 137 of FIG. 15. Upon searching down the hierarchy from the read-write inode 131,
3 it is noted that indirect blocks 133 and 135 and the data block 137 are “copy on write”
4 blocks. As shown in FIG. 16, new indirect blocks 139 and 140 and a new data block 141
5 have been allocated. The new data block 141 is a copy of the original data block 136
6 except that it includes the new data of the write operation. The new indirect block 140 is
7 a copy of the original indirect block 135 except it has a new pointer pointing to the new
8 data block 141 instead of the original data block 137. The new indirect block 139 is a
9 copy of the original indirect block 133 except it has a new pointer pointing to the new
10 indirect block 140 instead of the original indirect block 135. Also, the read-write inode
11 131 has been modified to replace the pointer to the original indirect block 133 with a
12 pointer to the new indirect block 139.

13 In some instances, a write to the read-write file will require the allocation of a
14 new data block without any copying from an original data block. This occurs when there
15 is a full block write, a partial block write to a hole in the file, or a partial block write to an
16 extended portion of a file. When there is a partial block write to a hole in the file or a
17 partial block write to the extended portion of a file, the partial block of new data is
18 written to the newly allocated data block, and the remaining portion of the newly
19 allocated data block is filled in with zero data.

20 It is possible that the UxFS layer will receive multiple concurrent writes that all
21 require new data to be written to the same newly allocated block. These multiple
22 concurrent writes need to be synchronized so that only one new block will be allocated
23 and the later one of the threads will not read old data from the original block and copy the

1 old data onto the new data from an earlier one of the threads. The UxFS layer detects the
2 first such write request and puts a corresponding entry into the partial block conflict
3 queue (73 in FIG. 4). The UxFS layer detects the second such write request, determines
4 that it is conflicting upon inspection of the partial block conflict queue, places an entry to
5 the second such write request in the partial write wait queue (74 in FIG. 4), and suspends
6 the write thread for the second such write request until the conflict is resolved.

7 FIG. 17 is a flowchart of steps in a write thread for performing the partial block
8 write operation of FIG. 13. In a first step 151 of FIG. 17, if the newly allocated file
9 system block (124 in FIG. 13) is not on the partial block conflict queue (73 in FIG. 4),
10 then execution branches to step 152. In step 152, the partial block write thread puts the
11 new block on the partial block conflict queue. In step 153, the partial block write thread
12 copies data that will not be overwritten by the partial block write, the data being copied
13 from the original file system block to the new file system block. In step 154,
14 asynchronous write operations are performed to write the new partial block of data to the
15 new block. In step 155, the partial block write thread gets the allocation mutex for the
16 file, commits the preallocated metadata (or the preallocated metadata is gathered and
17 committed upon servicing of the staging queue if a previous commit is in progress),
18 removes the new block from the partial block conflict queue, issues asynchronous writes
19 for any corresponding blocks on the partial write wait queue, and releases the allocation
20 mutex.

21 In step 151, if the newly allocated file system block was on the partial block
22 conflict queue, then execution continues to step 156. In step 156, the partial block write
23 thread puts a write callback on the partial write wait queue for the file. Then execution is

1 suspended until the callback occurs (from the completion of the asynchronous writes
2 issued in step 155). Upon resuming, in step 157, the partial block write thread gets the
3 allocation mutex for the file, commits the preallocated metadata (or the preallocated
4 metadata is gathered and committed upon servicing of the staging queue if a previous
5 commit is in progress), and releases the allocation mutex.

6 FIG. 18 shows steps in a write thread for allocating file blocks when writing to a
7 file having read-only versions. In a first step 161, if the file block being written to is not
8 shared with a read-only version, then execution branches to step 162 to write directly to
9 the block without any transaction. In other words, there is no need for allocating any
10 additional blocks.

11 In step 161, if the file block being written to is shared with a read-only version,
12 then execution continues to step 163. In step 163, if the file block being written to is an
13 indirect block, then execution branches to step 164. In step 164, a new indirect block is
14 allocated, the original indirect block content is copied to the new indirect block, and the
15 new metadata is written to the new indirect block synchronously. If the block's parent is
16 an indirect block shared with a read-only version, then a new indirect block is allocated
17 for copy-on-write of the new block pointer. Any other valid block pointers in this new
18 indirect block point to shared blocks, and therefore the most significant bit in each of
19 these other valid block pointers should be set (as indicated by the dotted line between the
20 indirect blocks 136 and 140 in FIG. 16). For example, just after the original indirect
21 block content is copied to the new indirect block, the most significant bit is set in all valid
22 block pointers in the new indirect block. As described above with respect to FIG. 16, this
23 copy-on-write may require one or more additional indirect blocks to be allocated (such as

1 indirect block 139 in FIG. 16). For example, the tree of a UxFS file may include up to
2 three levels of indirect blocks. All of the file blocks that need to be allocated can be
3 predetermined so that the allocation mutex for the file can be obtained, all of the new
4 blocks that are needed can be allocated together, and then the allocation mutex for the file
5 can be released.

6 In step 163, if the file block being written to is not an indirect block, then
7 execution continues to step 165. This is the case in which the file block being written to
8 is a data block. In step 165, if the write to the file block is not a partial write, then
9 execution branches to step 166. In step 166, a new data block is allocated and the block
10 of new data is written directly to the new data block. If the original block's parent is an
11 indirect block that is shared with a read-only version, then a new indirect block is
12 allocated for copy-on-write of the new block pointer. As described above with respect to
13 FIG. 16, this copy-on-write may require one or more additional indirect blocks to be
14 allocated.

15 In step 167, for the case of a partial write, execution continues from step 156 to
16 step 167 to use the partial write technique as described above with respect to FIG. 13 and
17 FIG. 17.

18 Various parts of the programming for handling a write thread the UxFS layer have
19 been described above with reference to FIGS. 11 to 18. Following is a listing of the steps
20 in the preferred implementation of this programming.

21 1. The write thread receives a write request specifying the source and
22 destination of the data to be written. The source is specified in terms of message buffers

1 and the message buffer header size. The destination is specified in terms of an offset and
2 number of bytes to be written.

3 2. The write thread calculates the starting and ending logical block number,
4 total block count, and determines whether the starting and ending blocks are partial
5 blocks.

6 3. The write thread gets the allocation mutex for the file.

7 4. The write thread searches the file tree along a path from the file inode to
8 the destination file blocks to determine whether there are any shared blocks along this
9 path. For each such shared block, a new data or indirect block is allocated
10 synchronously, as described above with reference to FIGS. 15, 16, and 18.

11 5. The write thread identifies partial blocks of write data using the starting
12 physical block number and the number of blocks to be written. Only the starting and
13 ending block to be written can be partial. Also, if some other thread got to these blocks
14 first, the block mapping may already exist and the “copy-on-write” will be done by the
15 prior thread. The partial block conflict queue is checked to determine whether such an
16 allocation and “copy-on-write” is being done by a prior thread. If so, the block write of
17 the present thread is added to the partial write wait queue, as described above with
18 reference to FIG. 17.

19 6. The write thread preallocates the metadata blocks.

20 7. The write thread releases the allocation mutex.

21 8. The write threads determine the state of the block write. The block write
22 can be in one of three states, namely:

- 1 1. Partial, in-progress writes. These are writes to blocks that are on the
2 conflict list. This write is deferred. The information to write out these
3 blocks is added to the partial write wait queue.
- 4 2. Whole Block Writes.
- 5 3. Partial, not-in-progress writes. These are partial writes to newly allocated
6 blocks, and are the first write to these blocks.
- 7 9. The I/O list is split apart if there are any non-contiguous areas to be
8 written.
- 9 10. Asynchronous write requests are issued for blocks in state 2 (full block
10 writes).
- 11 11. Synchronous read requests are issued for blocks in state 3 (Partial not-in-
12 progress writes).
- 13 12. Asynchronous write requests are issued for blocks in state 3.
- 14 13. The write thread waits for all writes to complete, including the ones in
15 state 1. The write thread waits for all asynchronous write callbacks. The asynchronous
16 writes for blocks in state 1 are actually issued by other threads.
- 17 14. The write thread gets the allocation mutex.
- 18 15. The write thread commits the preallocated metadata. The allocation lists
19 being committed are gathered together if a previous commit is in progress, and are
20 written out under the same logging lock as described above with reference to FIG. 12.
- 21 16. The write thread removes any blocks that the write thread had added to
22 partial block conflict queue, and issues asynchronous writes for corresponding blocks on
23 the partial write wait queue.

1 17. The write thread releases the allocation mutex. The write thread has
2 completed the write operation.

3
4 III. Maintenance of a file version set including read-only and read-write snapshot
5 copies of a production file.

6 As described above with reference to FIGS. 14 to 16 and 18, it is possible to use a
7 copy-on-write technique for creating a read-only snapshot of a Unix-based file. The
8 read-only snapshot can be used for non-disruptive backup by copying the read-only
9 snapshot to a backup media such as magnetic tape or optical disk. In this case the backup
10 is non-disruptive because the backup can be done as a background process while the
11 original read-write file can be accessed on a priority basis. Once a backup copy of the
12 read-only snapshot has been made, then the read-only snapshot can be deleted. For
13 example, the read-only snapshot is deleted by relinquishing the ownership of all of its
14 shared blocks back to the original read-write file, and then de-allocating all of the file
15 system blocks that are exclusively owned by the read-only snapshot.

16 Instead of using a single read-only snapshot for making a backup copy of a file, it
17 is possible to keep a series of read-only snapshots in the network file server. In this case,
18 when a crash occurs and the most recent snapshot is found to be corrupted, then an older
19 snapshot is immediately available for use in restoring the read-write file. Moreover, once
20 an entire copy of an initial snapshot has been migrated to the backup storage, only the
21 changes between the snapshots need be written to the backup storage in order to fully
22 recover all of the snapshots. In this case, there is a savings in backup processing time
23 and in backup storage capacity because more than one backup copy of each file system

1 block will neither be transmitted to the backup storage device nor stored in the backup
2 storage.

3 It is also desirable to provide a non-disruptive and virtually instantaneous
4 mechanism for making a read-write snapshot. For example, during the recovery process,
5 it is often desirable to create a temporary read-write copy of a read-only snapshot prior to
6 restoring the original read-write file after a system crash. Recovery can be attempted
7 upon the temporary read-write file, and then application programs can be tested upon the
8 temporary read-write copy. If a recovery program or an application program should
9 crash when using the temporary read-write copy, then the temporary read-write copy can
10 be deleted, and the recovery process can be restarted using another temporary read-write
11 copy of another read-only snapshot.

12 In order to facilitate the use of multiple read-only and read-write snapshot copies,
13 it is desirable to define a file version set including read-only and read-write snapshot
14 copies produced from an original read-write file. The original read-write file will be
15 referred to as the production file. The read-only snapshot copies will be referred to as
16 read-only versions, or simply versions. The read-write snapshot copies will be referred to
17 as branch files.

18 Shown in FIG. 19 is a preferred logical organization of such a file version set.
19 The file version set includes a production inode 171 for the production file, version
20 inodes 172, 173, 174 for a series of three read-only snapshots of the production file, and
21 two branch inodes 175, 176 for respective read-write copies of the most recent read-only
22 snapshot copy of the production file. The version set also includes a pool 177 of
23 exclusively owned and shared data blocks and indirect file blocks. Each data block or

1 indirect block in the pool 177 is linked to one or more of the inodes 171-176 either
2 directly or indirectly through an indirect block in the pool 177. As will be described
3 below with reference to FIGS. 21 to 23, the inodes 171 to 174 in the version set have a
4 modified format so that the inodes can be linked together via certain inode attributes.

5 Initially, the production file can contain a raw volume of allocated file blocks, or
6 the production file can be a sparse file that has no allocated blocks at creation time. For
7 the case of a sparse file, the initial read-only versions of the production file will be sparse
8 as well. As data is written to a sparse production file, the size of the file can grow up to a
9 pre-specified maximum number of blocks, and the maximum block size can then be
10 extended by moving the end-of-file (eof).

11 As shown in FIG. 20, a new production file is created as either a sparse file or a
12 fully preallocated file. For the case of a sparse file, execution branches from step 331 to
13 step 332 to initially allocate just the inode for the new sparse file. Otherwise, execution
14 continues from step 331 to step 333 to allocate an inode for the new fully preallocated
15 file. Then in step 334, all of the data blocks are allocated for a specified size for the new
16 fully preallocated file. Finally, in step 335, any and all indirect blocks are allocated for
17 the new fully preallocated file as needed to link any of the data blocks of the fully
18 preallocated file to the inode of the fully preallocated file. In other words, a fully
19 preallocated file is created with all of its metadata allocated, including all of its indirect
20 blocks and the data block pointers.

21 By initially allocating all of the metadata for a production file, the overhead
22 associated with the allocations, such as synchronization with concurrent allocations, is
23 eliminated for subsequent writes to the production file. A fully allocated production file

1 provides similar behavior as a storage volume, where all the data blocks are present at the
2 time of creation. A fully allocated production file, for example, is useful as a container
3 for storage objects that are known to be dense, such as video files or copies of raw disk.

4 The initial working file can also be created sparse by writing only to the inode and
5 last block of the file. The sparse file allows the production file to use only those blocks
6 that the client writes data to. This allows less disk blocks to be consumed initially. The
7 sparse file can then be used as the production file for the file version set. Since the new
8 production file after creating a snapshot copy uses new data blocks to write out the data,
9 it results in efficient data block usage, eliminating the need to allocate data blocks that
10 may never be used. The data block allocation scheme can allocate blocks for the new
11 working file in a way that can provide contiguity with the allocated blocks on the
12 previous snapshot copy allowing sequential access to the data blocks for better read
13 performance.

14 For management of the version set of FIG. 19, there is provided a protocol of
15 operations upon the version set. These operations include file creation, file deletion,
16 refresh, and recovery. File creation involves the creation of a read-only snapshot copy
17 from the production file or from a branch file, or the creation of a branch file off a read-
18 only version. File deletion involves the deletion of a read-only snapshot copy or a branch
19 file. Refresh involves discarding the contents of an existing read-only snapshot copy and
20 creating a new snapshot copy using the same name. Restore involves discarding the
21 contents of the production file and creating a new production file using the contents of a
22 specified read-only version.

1 FIG. 21 shows some of the fields of a conventional inode 180. The inode 180
2 includes a mode attribute (MODE) field 181, an access time attribute (ATIME) field 182,
3 an inode change time attribute (CTIME) field 183, one or more data block pointer fields
4 184, and one or more indirect block pointer fields 185.

5 FIG. 22 is a block diagram of an inode 190 in the file version set of FIG. 19. The
6 mode attribute 191 is set with a value IFVERSIONFILE indicating that the inode 190 is
7 for a file version set and the inode has a modified format, as further shown in FIG. 22.
8 The ATIME field 192 in the modified inode 190 stores a version pointer instead of an
9 access time. The CTIME field 193 in the modified inode 190 stores a branch pointer
10 instead of an inode change time. In addition to a data block pointer, the data block
11 pointer field 194 stores a non-owner flag 196 in the most significant bit position. The
12 non-owner flag 196 has a value of zero to indicate that the file is an owner of the data
13 block, and has a value of one to indicate that the file is a non-owner of the data block. In
14 addition to an indirect block pointer, the indirect block pointer field 195 stores a non-
15 owner flag 197 in the most significant bit position. The non-owner flag 197 has a value
16 of zero to indicate that the file is an owner of the indirect block, and has a value of one to
17 indicate that the file is a non-owner of the indirect block.

18 When there is only a production file, with no read-only snapshot copies, the
19 production file owns all of its blocks. When the first read-only snapshot copy file is
20 created, all of the blocks are passed to the new snapshot copy file and it becomes the
21 owner of all of the blocks. The production file still uses the same blocks and the same
22 blocks have identical contents (at least initially); however, it has become a non-owner of
23 those blocks. If any block of the production file is modified, then a new version of that

1 block is allocated and the production file will own that new block. (The new version of
2 the block will be a different block of storage mapped to the same logical address in the
3 file as the original version of the block.) As more snapshot files are created, different
4 snapshot files may own different versions of a block. The owner of any particular block
5 will always be the oldest snapshot copy that uses an identical version of a block, and the
6 oldest snapshot copy will always own all of its blocks. When a sparse file is used, each
7 time a new block is written to it will use the same UxFS allocation mechanism regardless
8 of who owns the data block, the production file or one of the snapshot copies.

9 The concept of a non-owner block is further extended, for indirect blocks, to
10 include the idea of a hierarchy of blocks. For indirect blocks and indirect block trees, if
11 the non-owner flag is set at any level of the tree, then the non-owner state is assumed for
12 all lower-level block pointers. For example, if a pointer to the first level indirect block is
13 marked as non-owner, then all of the data blocks that it points to are assumed to be non-
14 owner, regardless of the state of the non-owner flag in each of the individual block
15 pointer fields.

16 FIG. 23 further shows the use of the version pointers and the branch pointers for
17 linking the inodes 171-176 of the file version set introduced in FIG. 19. FIG. 23 shows
18 that the version pointers are used to form a linked list linking the production file inode
19 171 to all of the version inodes 172, 173, 174. Single links are used in the linked list to
20 conserve space within the conventional inode structure. The versions are linked from
21 most recent to least recent so that a new version inode can be created without modifying
22 other version inodes. The version pointer 201 of the production file 171 includes a most
23 significant bit that is set to indicate that the inode 171 is the inode of the production file.

1 The less significant bits of the version pointer 201 of the production file inode 171
2 contain the inode number of the most recent version if there is any read-only snapshot
3 copy in the version set, and if not, the inode number of the production file inode.

4 For example, in FIG. 23, the version pointer 201 of the production file inode 171
5 includes the inode number 16 of the third read-only snapshot copy inode 174. Each inode
6 172, 173, 174 of a read-only snapshot copy has a version pointer having a most
7 significant bit that is zero and an inode number of the inode of the next most recent read-
8 only version, or in the case of the oldest read-only version, the inode number of the inode
9 171 for the production file. The version pointer 204 of the inode 174 of the third version
10 contains the inode number 15 of the inode 173 of the second version. The version pointer
11 203 of the inode 174 of the second version contains the inode number 13 of the inode 172
12 of the first version. The version pointer 202 of the inode 172 of the first version contains
13 the inode number 10 of the production file inode 171.

14 The branch pointer in each inode has a most significant bit to that is set to indicate
15 the production file inode or a read-only version inode, and that is zero to indicate a
16 branch inode. The less significant bits of the branch pointer contain an inode number.
17 For the production file inode 171 or a read-only version inode 172, 173, 174, if the less
18 significant bits of the branch pointer contain the inode number of the inode, then there are
19 no branch files based on the production file or read-only snapshot copy file, respectively.
20 Otherwise, the less significant bits of the branch pointer in the production inode 171 or
21 version inode 172, 173, 174 include the inode number of the inode of the most recent
22 branch file based on the production file or read-only snapshot copy file, respectively.
23 The less significant bits of the branch pointer in a branch inode contain the inode number

1 of the next most recent branch file based on the same production file or read-only
2 snapshot copy file, or for the oldest branch inode, the inode number of the base
3 production or read-only snapshot copy file. In other words, if there are more than one
4 branch file based on the production file or a read-only version, then the branch pointers
5 are used to form a linked list of branch inodes off the base inode.

6 For example, in FIG. 23, the branch pointer 211 of the production file 171
7 contains the inode number 10 of the production file inode, since there are no branch files
8 based directly on the production file. The branch pointer 212 of the first version inode
9 172 contains the inode number 13 of the first version inode, since there are no branch
10 files based directly on the first read-only version. The branch pointer 213 of the second
11 version inode 173 contains the inode number 15 of the second version inode, since there
12 are no branch files based directly on the second read-only version. The branch pointer
13 214 of the third version inode 174 contains the inode number 18 of the second branch
14 inode 176. The branch pointer 216 of the second branch inode 176 contains the inode
15 number 17 of the first branch inode 175. The branch pointer 215 of the first branch inode
16 175 contains the inode number 16 of the third version inode 174.

17 In practice, it is desirable to prevent a user from creating a branch directly off the
18 production file, since otherwise it would not be possible to recover the branch file after a
19 disruption. The user can always create a read-write copy of the production file by first
20 creating a read-only snapshot copy of the production file and then creating a branch based
21 on the read-only version. If the branch file would be disrupted, then it could be recovered
22 from the read-only version.

1 Because the production file inode serves as an anchor for the snapshot chain, it is
2 desirable to prevent deletion of the production file if there are any snapshot files. The
3 snapshot files should be deleted first.

4 It is also desirable to prevent a read-only snapshot copy from being deleted if
5 there are any branch files based on the read-only version. Typically, any branch files
6 based on the read-only file would be deleted first. Instead of deleting a branch file, it
7 could be converted to a production file and unlinked from the base version, before
8 deletion of the base version. The branch file could be converted to a production file by a
9 background process of copying all blocks that are not owned by the branch file from the
10 base version to newly allocated blocks for the branch file. In the copying process, all of
11 non-owner flags would be cleared.

12 In addition, it is possible to write some changes to a branch file and then create a
13 read-only snapshot copy of the branch file. In this case, the version pointer in the branch
14 inode would contain the inode number of the inode of the read-only version of the branch
15 file. It would also be possible to create branches off this read-only version. In general,
16 the version inodes and the branch inodes could be linked in a hierarchy of version chains
17 and branch chains depending from the production inode 171.

18 Create and delete operations in a version set are synchronized. Further write
19 operations, which may allocate blocks, are synchronized with delete operations. A
20 shared global mutex (a version lock) is used to insure the integrity of the version and
21 branch chains while searching the chains for a file and while modifying the chains. To
22 prevent deadlocks, when concurrent locks are taken on more than one file in a chain, the
23 locking should be done from the head of the chain backwards through the chain. For

1 example, when two successive versions are concurrently locked to delete the earlier
2 version, a lock is first taken on the later version, and then a lock is taken on the earlier
3 version.

4 For each version set, only one create (snap, refresh, restore, etc.) or delete
5 operation may take place at a time. Additional create or delete operations are serialized,
6 because these operations may be changing more than one file in the version set. The
7 create operations are relatively quick and they will hold the global lock for the duration
8 of the operation. Delete operations can take significantly longer. Delete operations are
9 also controlled to prevent multiple delete operations from occurring at the same time.
10 For this purpose, a flag indicating that a delete operation is taking place and a condition
11 variable are maintained in the production file inode.

12 Typically, a Unix-based file system has a file check facility for checking the
13 integrity of the directories and linkages in a file system. This file check facility is
14 extended to recognize that a production file is in a file version set, and once a file version
15 set is found, to check the integrity of the branch and version chains, and to validate the
16 block pointers, the block ownership, and the block counts of the files in the version set.

17 FIG. 24 is a flowchart of a procedure for creating a read-only version of the
18 production file in the file version set of FIG. 19. In a first step 221, a new inode is
19 allocated for the read-only version. Then in step 222, the production file inode is locked.
20 In step 223, the production file inode is copied to the new inode for the version. In step
21 224, the new version inode is updated; for example, the version pointer is updated to link
22 the new version inode into the version chain off the production inode. In step 225, the
23 production file inode is updated; for example, the version pointer is updated to point to

1 the new version inode and the block pointer fields are updated (by setting the most
2 significant bits to set the non-owner flags) to show that the production file is a non-owner
3 of the pointed-to blocks. Then in step 226, the transaction is committed by writing an
4 entry for the new version creation to the log, and writing the production file inode and the
5 new version inode to the file system in storage. Finally, in step 227, the lock on the
6 production file inode is released.

7 FIG. 25 is a flowchart of a procedure for creating a read-write branch off a base
8 version in the file version set of FIG. 19. In a first step 231, a new inode is allocated for
9 the read-write branch. In step 232, the base version inode is locked. Then in step 233,
10 the base version inode is copied to the new inode for the branch. In step 234, the new
11 branch inode is updated; for example, the branch pointer is set to link the new branch
12 inode into the branch chain off the base inode, and the block pointer fields are updated
13 (by setting the non-owner flags in the most significant bits) to indicate that the branch file
14 is a non-owner. In step 235, the base version inode is updated; for example, the branch
15 pointer is set to point to the new branch inode. In step 236, the transaction is committed;
16 for example, by writing an entry into the log indicating the creation of the new read-write
17 branch off the base version, and by writing the new branch inode and the updated base
18 inode to the file system in storage. Finally, in step 237, the lock on the base version
19 inode is released.

20 FIG. 26 shows a procedure for deleting a read-only version in the file version set
21 of FIG. 19, while retaining the next most recent snapshot copy (or the production file,
22 when the snapshot copy being deleted is the most recent read-only version). This
23 involves deleting blocks that are exclusively owned by the snapshot copy being deleted,

1 and retaining blocks that are shared between the snapshot copy being deleted and the next
2 most recent version.

3 In a first step 241 of FIG. 26, a lock is taken on the inode of the read-only
4 snapshot copy and the inode of the next most recent snapshot copy (or the production file
5 if the read-only snapshot copy being delete is the most recent read-only version). The
6 lock prevents the deletion operation from changing the file mapping at the same time that
7 new allocations are being done. If the read-only snapshot copy being deleted is the most
8 recent read-only version, then this lock on the production file is taken in shared mode by
9 writes (and allocations) to prevent blocks owned by the most recent version and not
10 owned by the production file from being passed up to the production file (in step 242) at
11 the same time that new blocks are being allocated.

12 In step 242, there is begun a search for blocks indexed in the inode of the read-
13 only version and corresponding blocks in the inode of the next most recent version (or in
14 the production file if the read-only snapshot copy being deleted is the most recent read-
15 only version). A block in the next most recent version corresponds to a block in the read-
16 only snapshot copy being deleted if the two blocks map to the same range of logical
17 addresses in the two files. The corresponding block may be an identical version of a
18 block (i.e., the same block of storage), in which case the contents will also be the same
19 (because the copy-on-write technique would be used to allocate a new storage block if
20 the contents would change).

21 The search for the corresponding blocks is referred to as a coalescing pass. The
22 objective is to locate blocks that are exclusively owned by the read-only snapshot copy so
23 that these blocks can be freed. Another objective is to locate shared blocks that are

1 owned by the read-only snapshot copy so that ownership of these blocks can be passed to
2 the next read-only snapshot copy (or the production file if the read-only version being
3 deleted is the most recent read-only version). To carry out these objectives, in step 243,
4 the ownership state of each block in the version being deleted is inspected, and a
5 corresponding action is taken depending on the ownership of the block. If the block is
6 not owned by the version being deleted, then an identical version of the block is shared
7 with and owned by an earlier snapshot copy. Also, by inheritance, all of the descendants
8 of the block in the block hierarchy are shared with and owned by an earlier snapshot
9 copy. Therefore, the block (and all of its descendants) can be ignored. The searching
10 process skips over the block and all of its descendants.

11 If the block is owned by the snapshot copy being deleted, then an action is taken
12 depending on the state of the corresponding block in the next most recent version (or the
13 production file if the read-only snapshot copy being deleted is the latest version). If the
14 corresponding block in the next most recent version is not owned, then an identical
15 version of the block is shared between the read-only version being deleted and the next
16 most recent version, and ownership of the block is passed from the read-only version
17 being deleted and the next most recent version. As blocks are passed, the block count is
18 incremented for the next most recent version. If the block being passed is an indirect
19 block, then its descendants become passed by inheritance. However, the indirect block
20 (and any indirect block descendants) should be traversed to count the number of
21 descendants in order to increment the block count for the next most recent snapshot copy
22 by the number of descendants. A function is provided to do the counting for one indirect

1 block, and this function may be called recursively for second and third level indirect
2 trees.

3 If the block is owned by the read-only file version being deleted and the
4 corresponding block in the next most recent file snapshot is owned by the next most
5 recent version, then the block was modified between the read-only snapshot being deleted
6 and the next most recent version. In this case, the read-only snapshot copy being deleted
7 has exclusive ownership of its version of the block, and its version of the block can be
8 freed. If the block is not found in the next most recent version (for example because the
9 extent of the file had been reduced), then the read-only snapshot copy being deleted has
10 exclusive ownership of the block, and the block can be freed.

11 Finally, in step 244, when the search for blocks has been completed, the inode of
12 the read-only snapshot copy being deleted is deallocated, and the lock is released.

13 The deletion of blocks from the read-only snapshot copy being deleted can be
14 done in such a way that truncation occurs from the end of the file backwards. In this
15 case, the file size can be used as a processing indicator, and the deletion process can be
16 halted and restarted. The coalescing and cleanup of the file can be done asynchronously,
17 although only one file deletion from the version set will be performed at any given time.

18 In a preferred implementation, the coalescing and cleanup of a file is done by a
19 program loop that executes a series of transactions. Each pass through the program loop
20 executes one transaction. Each transaction is logged, so the coalescing and cleanup can
21 be resumed if interrupted by a system crash. During each transaction, an exclusive lock
22 is held on the next most recent version (or the production file, if the snapshot copy being
23 deleted is the most recent read-only version). This prevents any attempt to allocate

1 blocks in the locked file. The exclusive lock is released at the end of processing for each
2 transaction, in order for any conflicting processes to make forward progress.

3 The process of deleting versions can be simplified when all of the files in the
4 version set are deleted. In this case, all of the blocks in the version set are deallocated. In
5 addition, the deletion of multiple successive versions can be optimized. Only a single
6 coalescing pass is needed to pass blocks that are owned by the successive versions being
7 deleted but shared with the next most recent version being retained. There is no need to
8 pass blocks between two successive versions that will both be deleted.

9 There is no need for passing blocks when a branch file is deleted. Any blocks that
10 are owned by the branch are deallocated, and any non-owner blocks are ignored.

11 FIGS. 27-28 show details of the persistent reservation mechanism ensuring that
12 the sharing of file blocks among the files in the version set of FIG. 19 is not likely to
13 result in a shortage of file blocks when writing to the production file or a branch file. In a
14 first step 251, a number of free blocks are reserved for each read-write file. The number
15 is maintained as a "block reservation" attribute for the read-write file. In step 252, the
16 number of blocks in each file is maintained as a "block count" attribute for the read-write
17 file. In step 253, when a new block is allocated to the read-write file from the block
18 reservation for the file, the block reservation is decremented, and the block count is
19 incremented. Additional free blocks are reserved to prevent the block reservation from
20 becoming negative, or else the allocation fails. In step 254, when a block is removed
21 from the read-write file, then an additional block can be reserved for the file. In this case,
22 the block reservation for the file is incremented, and the block count for the file is
23 decremented.

1 Continuing in step 255 of FIG. 28, the creation of a read-only snapshot copy of a
2 read-write file will fail unless there can be reserved a number of free blocks equal to the
3 block count of the read-write file. For example, more free blocks are reserved as the
4 block reservation count of the read-write file is incremented by the number of blocks that
5 become shared with the new read-only file.

6 In step 256, the creation of a read-write branch of a read-only base snapshot copy
7 will fail unless there can be reserved a number of free blocks equal to the block count of
8 the read-only base version. For example, more free blocks are reserved as the block
9 reservation of the branch file is incremented by the number of blocks that become shared
10 with the new read-write branch file.

11 In step 267, a restore of the production file with a read-only snapshot copy will
12 fail if the block count of the read-only snapshot copy exceeds the block count of the
13 production file unless there can be reserved a number of free blocks equal to the
14 difference between the block count of the read-only snapshot copy and the block count of
15 the production file. For example, the block reservation of the production file is
16 incremented by the original block count of the production file, decremented by the block
17 count of the read-only version, and any deficiency is made up by incrementing the block
18 reservation as additional free blocks are reserved for the production file.

19 FIG. 29 is a state diagram for the process of restoring a production file with a
20 read-only version. The state diagram has an initial state 261 of the original production
21 file, an intermediate state 262 in which the version set has been prepared for a restore
22 operation, and a final state 263 in which the production file has been restored. The
23 process of restoring the production file is provided with a distinct intermediate state

1 because it is possible that the restore operation may fail or it may be desirable to provide
2 the user with an option to abort the restoration process, for example, because sufficient
3 free file system blocks are not available to satisfy the persistent reservation requirement.
4 Therefore, the restoration process has been configured for a two-phase commit process,
5 in which the first phase is to prepare for a restore operation, and the second phase is to
6 either abort the restore operation or commit the restore operation.

7 Once a process capable of failure has been configured for such a two-phase
8 commit process, then it can be used in the well-known two-phase distributed commitment
9 protocol. In the two-phase distributed commitment protocol, the preparation and
10 commitment can be done at multiple distributed sites under management of a single
11 controller. In the first phase, the preparation at all of the sites is performed at the request
12 of the controller, and the results are reported back to the controller. If all sites report
13 back that the preparation has been successful, then the controller may request all of the
14 sites to commit to completing the process. In this case, it is highly probable that the
15 process will be completed everywhere. However, if any one of the sites reports back that
16 its preparation has been unsuccessful, then the controller may request all of the sites to
17 abort their preparation.

18 For restoring files, the two-phase distributed commitment protocol could be
19 useful for preparing to restore multiple files in a distributed data storage system. The
20 files could be distributed across a network and stored in different network file servers. If
21 the preparation for restoration of all of the files would be successful, then the restoration
22 of all of the files would be committed. If the preparation for restoration of any of the
23 files would be unsuccessful, then the restoration of all of the files would be aborted. The

1 preparation for the restoration process could ensure, to a high probability, that all of the
2 files in the file system could be restored together, or else none of them would be restored.

3 FIG. 30 shows a procedure for preparing for the restoration of the production file.
4 In a first step 271, a branch file copy is created from a specified base version. The base
5 version is the read-only snapshot copy to be used for restoring the production file. Also
6 an attempt is made to reserve the difference between the block count of the specified base
7 version and the block count of the production file. In step 272, if the restoration has been
8 prepared, then execution returns reporting success. Otherwise, execution returns
9 reporting failure. For example, execution could return with a fatal error if the specified
10 base version has been corrupted so that no branch file copy could be created. Execution
11 could also return with an indication that creation of the branch file copy was successful
12 but there were insufficient resources for persistent reservation.

13 FIG. 31 shows a procedure for aborting the restoration of the production file. In
14 step 281, the new branch file (created during preparation for the restore) is discarded.
15 Read-write access may continue with the original production file.

16 FIG. 32 shows a procedure for committing the restoration of the production file.
17 In step 291, the new branch file (created during the preparation for the restore) assumes
18 the identity of the production file. This is done by unlinking the branch file inode from
19 the branch chain off the base version inode, linking the branch file inode into the version
20 chain in lieu of the production file inode, and changing the pointer in the parent directory
21 of the production file to point to the branch file inode in lieu of the production file inode.
22 Then the old production file inode and the blocks owned by the old production file are
23 deallocated. Unless a nondestructive restore option is selected, any read-only versions

1 more recent than the base version are also deleted by deallocating all of their owned
2 blocks and then deallocating their inodes.

3 A refresh of a read-only snapshot copy takes an existing version file, discards it
4 contents, and creates a new version for the snapshot file. The new version is a snapshot
5 copy of the present state of the production file. FIG. 33 shows a preferred procedure. In
6 step 301, a new inode is created, and the contents of the original version inode are copied
7 into the new inode. In step 302, the new inode is linked into the version chain in lieu of
8 the original version inode. In step 303, the original version inode is used to create a new
9 snapshot of the production file. In other words, the production file inode is copied to the
10 original version inode, the original version inode is linked into the version chain as the
11 most recent version, and the non-owner flags are set in the production file inode. Then in
12 step 304, the old read-only snapshot copy of the new inode is scheduled for asynchronous
13 deletion. In this fashion, the refreshed snapshot can become available for user access
14 before the old snapshot copy is deleted.

15 It is desirable to provide users with a convenient method of referencing the
16 various files in a version set. A preferred method is to provide a hierarchical naming
17 convention similar to a hierarchical path name common for Unix-based file; for example,
18 a path name for a Unix-based file is typically in the form of:

19 DirectoryName\SubDirectoryName\...\FileName.

20 For referencing files in a version set, a suitable hierarchical naming convention is in the
21 form of:

22 ProductionFileName [\$VersionName][\$\$BranchName]....

1 In other words, a single occurrence of the “\$” symbol is used as a delimiter to indicate a
2 following version name, and a double occurrence of the “\$\$” symbol is used as a
3 delimiter to indicate a following branch name. In accordance with this convention, the
4 six files in the version set of FIG. 19 could have the following file names:

5
6 Production Inode 171: ProductionFileName
7 Version 1 Inode 172: ProductionFileName\$1
8 Version 2 Inode 173: ProductionFileName\$2
9 Version 1 Inode 174: ProductionFileName\$3
10 Branch 1 Inode 175: ProductionFileName\$3\$1
11 Branch 2 Inode 176: ProductionFileName\$3\$2

12
13 This naming convention would have the advantage that all of the files in the
14 version set could share the same NFS file handle or CIFS file id. In addition, the naming
15 convention would have the advantage that a file name matching the pattern could trigger
16 the creation of a new snapshot copy or branch file. For example, if a request to create a
17 new version specified an existing production file name followed by the delimiter “\$”
18 followed by a version name that did not exist, then a new snapshot of the production file
19 would be created having the specified version name. The file handle returned would be
20 that of the production file.

21 FIGS. 34 and 35 show a procedure for parsing a file name in accordance with the
22 above convention. In a first step 311, version chain scanning is set to begin at the
23 production inode. Then in step 312, the production file name is parsed from the name of

1 the file in the version set. In step 313, if an end of input is reached in the parsing of name
2 of the file in the version set, then execution returns indicating that the production file is to
3 be accessed. Otherwise, execution continues to step 314 to get the next character from
4 the file name input. In step 315, if this next character is not the "\$" character, then
5 execution returns reporting a format error. Otherwise, execution continues to step 316, to
6 parse a version name X and scan the version chain until the inode is found for the version
7 named X. In step 317, if an end of input is reached in the parsing of the version number,
8 then execution returns indicating that the read-only version X of the production file
9 system is to be accessed. Otherwise, execution continues to step 318 in FIG. 35.

10 In step 318 of FIG. 35, the next two characters are obtained from the input of the
11 name of the file in the version set. In step 319, the next two characters are not "\$\$", then
12 execution returns reporting a format error. Otherwise, execution continues to step 320 to
13 parse a branch name Y and scan the branch chain off the version named X until the
14 branch named Y is found. In step 321, if an end of input of the name of the file in the
15 version set has been reached, then execution returns indicating that the branch Y off the
16 read-only snapshot copy X is to be accessed. Otherwise, execution continues to step 322.
17 In step 322, the next character is obtained from the input of the name of the file in the
18 version set. In step 323, if the next character is not "\$", then execution returns reporting
19 a format error. Otherwise, execution continues to step 324 to set the version chain
20 scanning to begin at the inode of branch Y of version X. After step 324, execution loops
21 back to step 316 of FIG. 34.

22 An alternative naming convention could use a directory for the version set. The
23 directory could have an entry for each file in the version set, and an arbitrary name could

1 be assigned to each file in the version set. The directory for the version set could provide
2 a means for locating a branch file that would become unlinked from its base version
3 when its base version is deleted, or locating versions that might be retained after deletion
4 of the production file. This alternative, however, involves additional processing overhead
5 for maintaining the directory entries and keeping track of the directory itself.

6 Another alternative is to use pseudo directories. Each pseudo directory could
7 have a version date or user supplied label associated with it. Also, it could have a
8 specific file system version level number. Reading the pseudo directory could return a
9 list the files that had a version number less than or equal to the version number of the
10 pseudo directory. This has an advantage in that it is somewhat easier to manage older file
11 versions, since they are collected together in the pseudo directories. This alternative
12 would require the production file to exist as a name anchor and would also involve
13 additional processing time for maintaining the pseudo directories.

14 In view of the above, there have been described a multi-threaded write interface
15 for increasing the single file write throughput of a file server. The write interface allows
16 multiple concurrent writes to the same file and handles metadata updates using sector
17 level locking. The write interface provides permission management to access the data
18 blocks of the file in parallel, ensures correct use and update of indirect blocks in the tree
19 of the file, preallocates file blocks when the file is extended, and solves access conflicts
20 for concurrent reads and writes to the same block, and permits the use of pipelined
21 processors. The write interface preallocates file metadata to prevent multiple writers
22 from allocating the same block. For example, a write operation includes obtaining a per
23 file allocation mutex (mutually exclusive lock), reserving a metadata block, releasing the

1 allocation mutex, issuing an asynchronous write request for writing to the file, waiting for
2 the asynchronous write request to complete, obtaining the allocation mutex, committing
3 the preallocated metadata block, and releasing the allocation mutex. Since no locks are
4 held during the writing of data to the on-disk storage and this data write takes the
5 majority of the time, the method enhances concurrency while maintaining data integrity.

6